

# Image analysis for the identification of coherent structures in plasma

Nicole S. Love and Chandrika Kamath

Center for Applied Scientific Computing  
Lawrence Livermore National Laboratory  
7000 East Ave., Livermore, USA

## ABSTRACT

Turbulence at the edge of the plasma in a nuclear fusion reactor can cause loss of confinement of the plasma. In an effort to study the edge turbulence, the National Spherical Torus Experiment uses a gas puff imaging (GPI) diagnostic to capture images of the turbulence. A gas puff is injected into the torus and visible light emission from the gas cloud is captured by an ultra high-speed camera. Our goal is to detect and track coherent structures in the GPI images to improve our understanding of plasma edge turbulence. In this paper, we present results from various segmentation methods for the identification of the coherent structures. We consider three categories of methods - immersion-based, region-growing, and model-based - and empirically evaluate their performance on four sample sequences. Our preliminary results indicate that while some methods can be sensitive to the settings of parameters, others show promise in being able to detect the coherent structures.

**Keywords:** segmentation, object detection, coherent structures, plasma edge turbulence

## 1. INTRODUCTION

The success of magnetically-confined fusion reactors, such as ITER,<sup>1</sup> is dependent on two key factors - first, the plasma has to be hot enough so that the particles can fuse, and second, it has to be confined long enough so that the particles do fuse. This can be challenging for many reasons, one of which is the presence of fine-scale turbulence which causes leakage of plasma and particles from the center of the reactor to the edge. This leakage could result in a significant heat loss from the plasma, loss of confinement of the particles, as well as erosion or vaporization of the containment wall of the reactor.

To understand the physical mechanisms behind this fine-scale turbulence, scientists have exploited 2-D optical imaging technology to obtain images of the edge turbulence in experimental reactors. These images have indicated the presence of “coherent” structures which “retain their geometry over many characteristic lengths of motion”.<sup>2</sup> By characterizing these structures and tracking their motion over time, physicists can improve their understanding of edge turbulence and its role in the transport of heat and particles to the edge of the plasma.

In this paper, we describe the use of image analysis techniques to identify the coherent structures in experimental images of plasma turbulence. In Section 2, we first describe the problem, the data, and the challenges we face in the analysis. Next, in Section 3 we discuss various techniques to identify the structures in the images. Section 4 presents empirical results for different image sequences. Finally, we conclude with a summary in Section 5.

## 2. PROBLEM DESCRIPTION

We conduct our analysis of edge turbulence using images from the National Spherical Torus Experiment (NSTX) at the Princeton Plasma Physics Laboratory. NSTX is a magnetic fusion device in the form of a spherical torus (that is, a low aspect ratio tokamak) with several state-of-the-art diagnostics. One of these is the gas puff imaging (GPI) optical diagnostic<sup>3</sup> which is used to view the coherent structures in edge turbulence. The turbulence is known to consist of density “filaments” which are highly elongated in the direction of the local magnetic field. To view the 2-D structure perpendicular to the magnetic field, a neutral deuterium gas cloud is injected into the torus and the visible light from the intersection of the gas cloud with the filament is captured as a “blob” using a camera which views the filament along the direction of the magnetic field. The term “blob” refers to the

shape of the 3-D filament seen in the 2-D image. It is created with concentrations of density, temperature, and vorticity higher than the surrounding plasma.<sup>4</sup>

Figure 1, first row, shows a frame each from four different image sequences from NSTX. These sequences are obtained using an ultra-high speed, high resolution, PSI-5 camera, with each sequence consisting of 300 frames taken at 250,000 frames/second. Each 16-bit image is  $64 \times 64$  pixels. The original raw images are often noisy, with isolated bright and dark pixels making it difficult to see the blobs in the data. The images are first processed using a 2-D  $3 \times 3$  median filter to remove the noise spikes, followed by further smoothing using an  $11 \times 11$  Gaussian filter with  $\sigma = 0.4$ . The resulting de-noised images are shown in Figure 1, second row.

Further processing is required before we can identify the extent of the blobs in these images. The background or quiescent intensity must be removed from the images. This is the intensity which would have been observed in the absence of the blobs. For short sequences, we can represent this quiescent emission profile by using either the mean or the median of the 300 frames. This can be either subtracted out or divided out from the de-noised images. For the sequences considered, all four options produce similar results. Figure 1, third row shows the quiescent image obtained by taking the median, and the fourth row shows the final result, after the background has been divided out from the de-noised image. As expected by the physicists, the quiescent image has roughly constant intensity vertically, but varying intensity horizontally. The analysis in the remainder of the paper is done using the images from the fourth row of Figure 1.

Figure 2 shows the movement of the blobs over time using frames 10 through 19 from the sequence 113734. In these images, the edge of the plasma is to the right of the image. These images show that the blob behavior over time can be quite complex, with a blob initially moving downward and merging with a fainter blob which becomes brighter and splits off on its own, eventually becoming quite faint.

Our main motivation for identifying, extracting, characterizing, and tracking the coherent structures is to compare the experimental data with theory so that we can validate or refine the theory. A key challenge in the analysis is that the coherent structures are poorly understood empirically and not understood theoretically. Consequently, the physics of these structures is under intensive investigation.<sup>2,5,6</sup> From the analysis viewpoint, this unfortunately implies that we cannot compare the effectiveness of our blob-extraction methods by comparing our results with, say, a known ground-truth image, as there is no such image. What makes the problem even harder is that we cannot let our current understanding of the theory influence the results as the goal of the analysis is to validate the theory.

In the absence of ground-truth images, we need to consider alternative ways of evaluating different blob-extraction methods. For example, if an algorithm depends on many parameters making it difficult to select the best set of parameters, or the results of an algorithm are very sensitive to the values of the parameters, then it is unlikely to be a good candidate for extracting the blobs. This is because it would be difficult to ensure that the results are a true reflection of the data and not an artifact of the choice of algorithms or the parameters used. We discuss our approach to evaluation of the different methods further in Section 4.

### 3. IMAGE SEGMENTATION

To identify and extract the coherent structures in the images of edge plasma, we use image segmentation methods.<sup>7</sup> Image segmentation is the process of partitioning an image into regions whose pixels are similar to each other or whose pixels are different from the pixels of other neighboring regions. Our focus is on techniques which are simple and have relatively few parameters. In this study, we do not exploit knowledge of the structures in neighboring frames, though coherent structure have been identified using their motion.<sup>8</sup>

We next describe the three segmentation methods we considered, namely, immersion-based methods, region growing, and model-based methods. These are all variants of region-growing approaches. The diffuse boundary of the blobs makes it difficult to use gradient-based methods such as the Canny edge detector.

#### 3.1 Immersion-Based Methods

Immersion-based techniques consider the image as a topographic relief map, where elevations in the topography correspond to intensities in the images. The term “immersion” reflects the fact that segmentation is done by

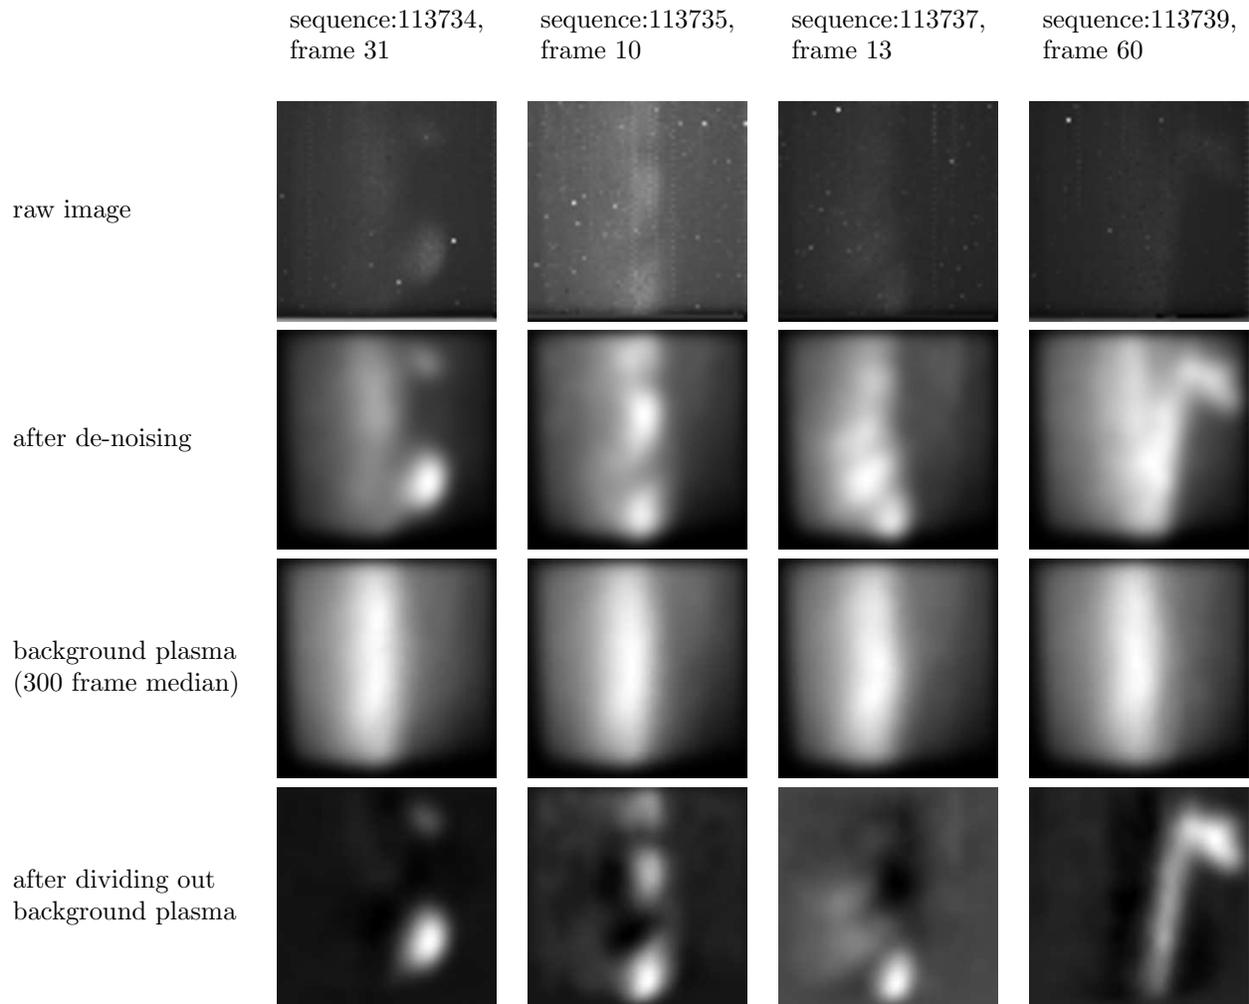


Figure 1. Removal of noise and background plasma from the raw image data for 4 frames from 4 different sequences. First row: original images. Second row: after filtering to remove noise. Third row: the quiescent or background plasma, defined as the median of the 300 frames of each sequence. Fourth row: Filtered image after the quiescent plasma intensity has been divided out.

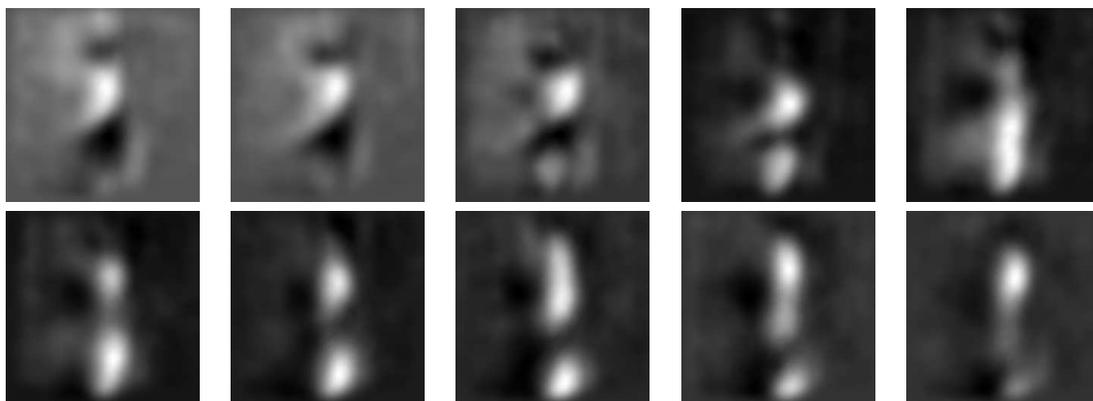


Figure 2. An example of 10 consecutive frames (10-19) from sequence 113734 after dividing out the background demonstrating the movement of plasma towards the edge (right).

simulating the immersion of the topographic relief map in water. Imagine there is a hole at each local minimum (or valley) of the image. When this relief map is immersed in water, it starts to fill from these holes to the extent determined by the depth of immersion. The immersion process is a form of thresholding where the depth of immersion determines the threshold. We examined several immersion-based methods as described next.

### 3.1.1 The basic immersion method

In the straight-forward implementation of the immersion method, the image is immersed to a specific depth,  $d$ . Full immersion occurs when all the pixels are immersed. Immersed pixels that are connected are considered to be a foreground region forming the coherent structure of interest. The depth is specified as a fraction of the intensity range of the image. Since our interest is in the higher intensity regions, the images are inverted prior to the application of immersion.

Figure 3 shows our four sample images after segmentation with three different immersion depths. The immersion at depth 0.25 captures the high intensity regions well. However, the structures with a low intensity are not captured at a depth of 0.25 or 0.375. Even when the depth is increased to 0.5, the faint region at the top of sequence 113734 frame 31 is not detected; a further increase is needed to detect this blob.

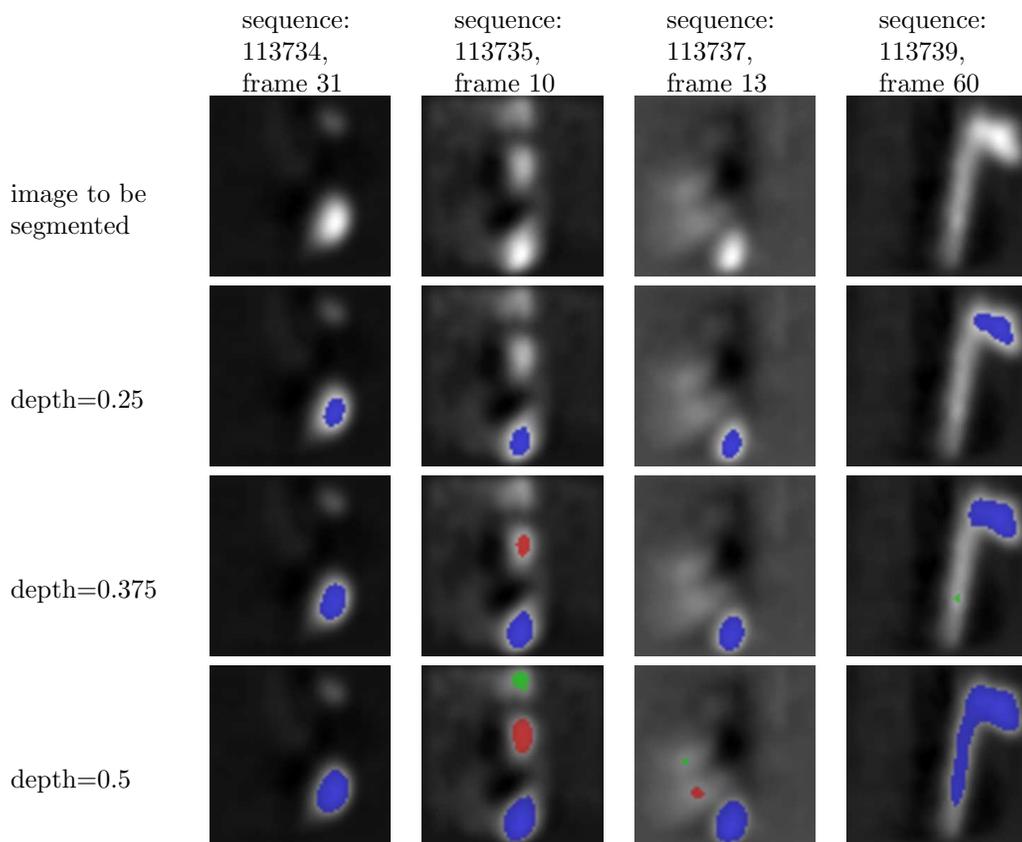


Figure 3. Results from the basic immersion method: Row one - image to be segmented; row two - immersion depth of 0.25; row three - immersion depth of 0.375; row four - immersion depth of 0.5.

### 3.1.2 The watershed method

The watershed algorithm addresses some of the drawbacks of the basic immersion method by identifying watershed lines which divide the valleys in the image. There are many possible implementations due to the ambiguity of the position of the watershed lines. We chose the algorithm of Vincent and Soille,<sup>9</sup> which is designed for fast computation and produces thin watershed lines. It is applied to the gradient image to help distinguish between the hills and the valleys. The algorithm proceeds by immersion until all pixels are labeled as a region or divider of a region (watershed lines). A region is defined as the area which is filled from the same local minimum.

The watershed algorithm is known to over-segment an image due to variations in intensities within a region. This can be clearly seen in Figure 4, row 2 which shows the results obtained by segmenting our sample images using the Vincent-Soille algorithm. We address this problem using the following two techniques:

- **Constrained watershed algorithm:** In this algorithm,<sup>10</sup> the image is first preprocessed to identify foreground and background marker regions. The watershed algorithm is then applied to the gradient of the image, and the markers are used to constrain the algorithm. The process is summarized as follows:
  - *Clean the image to simplify identification of foreground markers:* First, small blemishes in the image are removed by using the morphological operations<sup>7</sup> of open-by-reconstruction followed by close-by-reconstruction. This is an optional step which does not affect the overall shapes of the objects but makes it easier to identify flat regions which are at the maxima of each object in the image. Without this step, for our data, these maxima are often just a single pixel in size.
  - *Identify foreground markers:* Next, regional maxima are obtained by considering connected pixels with the same intensity value surrounded by boundary pixels which are strictly lower in intensity.
  - *Identify background markers:* The background markers are the pixels determined to be part of the background. They are obtained by thresholding the image. The threshold is the bin boundary of the image histogram which maximizes the inter-class variance:

$$\sigma^2 = w_0 w_1 (m_0 - m_1)^2,$$

where  $w_i$  and  $m_i$  are the probability and mean, respectively, of class  $i$ .<sup>11</sup> The two classes are the foreground coherent structures and the background.

- *Set new minima:* To obtain watershed lines between the foreground and the background markers, the markers are forced to be the minima of the gradient image by setting the intensity at each marker pixel to be  $-\infty$ .
- *Apply the watershed algorithm:* The Vincent and Soille watershed algorithm is then applied to the gradient image with the new minima.

Figure 4 shows the foreground and background markers for our sample images. Omitting the cleanup step results in several very small foreground regions (row 3). With the cleanup step, some of the regions increase in size and coalesce, while the less significant ones disappear altogether (row 4). Row 5 shows the background markers. Note that due to the definition of the foreground markers, some of them in the low intensity regions overlap with the background markers.

- **Watershed merging**

The second approach to reducing over-segmentation is based on post-processing. The watershed-merging algorithm<sup>12</sup> merges regions based on similarity using the following metric:

$$\delta(R_i, R_j) = (\mu(R_i) - \mu(R_j))^2 a(i, j),$$

where  $\mu(R_i)$  is the mean of region  $i$  and  $a(i, j)$  is 1 if region  $i$  and region  $j$  are adjacent and  $+\infty$  otherwise. A large value of  $\delta(R_i, R_j)$  indicates that the regions  $R_i$  and  $R_j$  are very dis-similar, so  $\delta$  acts as a dissimilarity metric. As an initial step, we merge all regions in the watershed which are completely contained in the background markers. The background region, along with the remaining regions, which have at least one pixel not in the background, are considered for further processing with subsequent merges based on descending order of similarity. Merging continues until a threshold,  $t = w \max(\delta(R_i, R_j))$ , is reached, where the maximum value is obtained from the dissimilarity of the initial regions. In our work, we use  $w = 0.02$ . Note that each time two regions are merged, the mean of the new region must be calculated.

Figure 4 shows that both the constrained watershed method (row 6) and the watershed merging approach (row 7) significantly reduce the over-segmentation present in the standard watershed technique (row 2). The regions shown are the ones which contain at least one foreground marker pixel which is simultaneously not a background marker pixel.

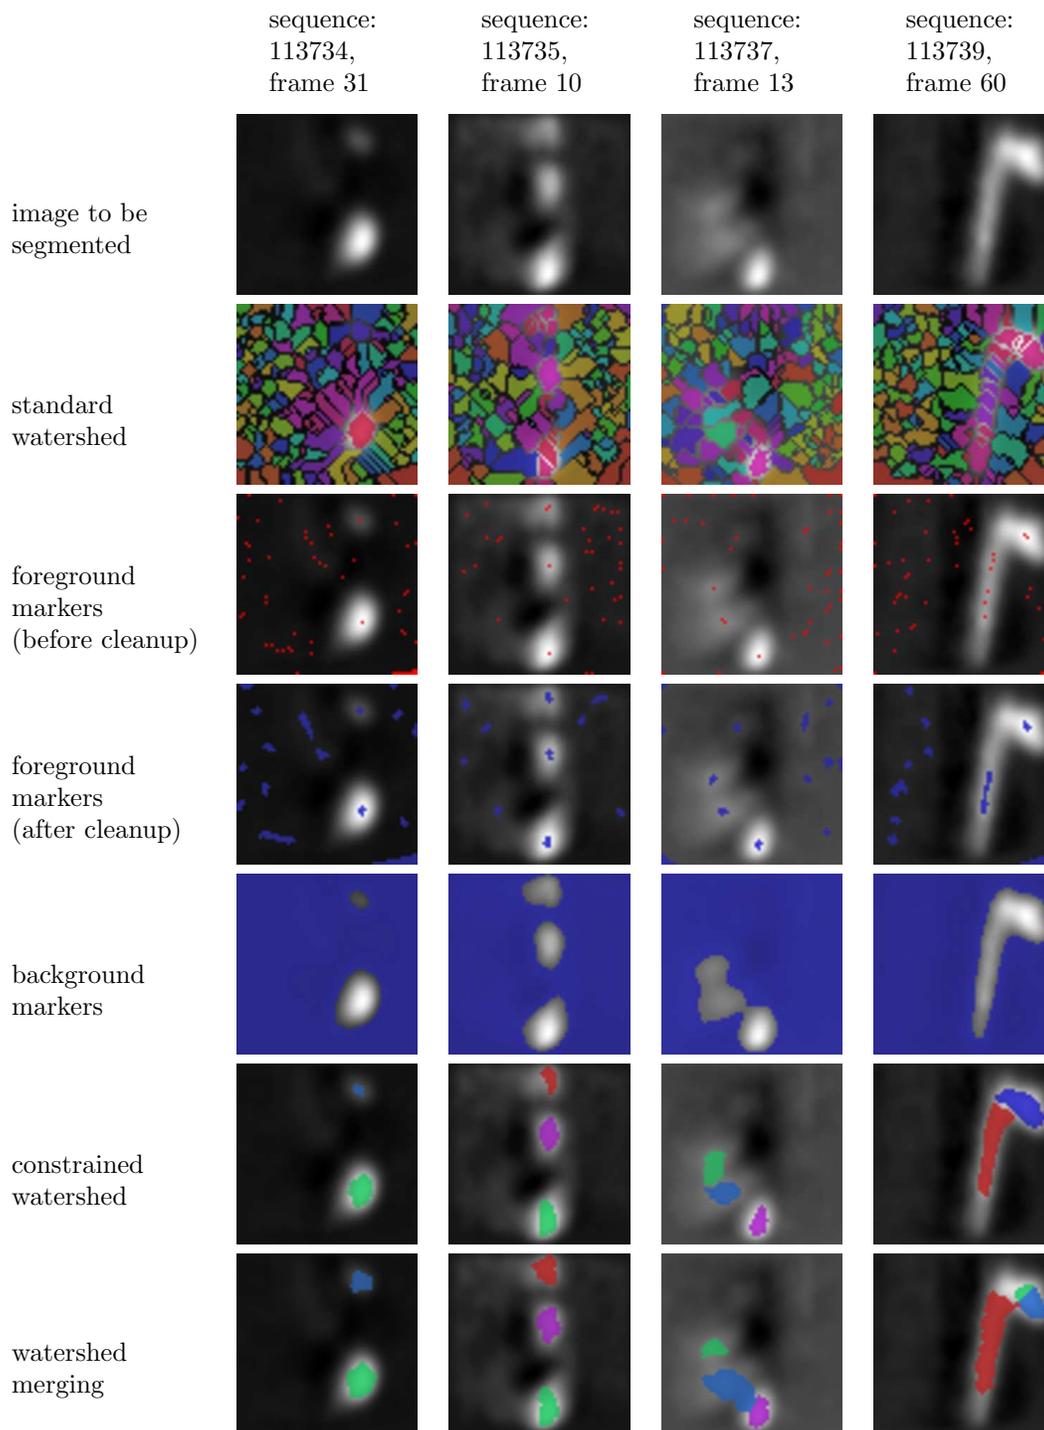


Figure 4. Results for watershed segmentation. Row 1 - image to be segmented; row 2 - standard watershed segmentation indicating the over-segmentation of the image; row 3 - foreground markers obtained without the cleanup step; row 4 - foreground markers obtained after the cleanup step; row 5 - background markers; row 6 - constrained watershed segmentation; row 7 - results with watershed merging

## 3.2 Region Growing Methods

The region growing methods for image segmentation start from small regions called “seeds” which are grown in a constrained manner. In our work, we choose the seeds to be the pixels from the foreground and background marker regions described in Section 3.1.2. We consider two region growing methods, one based on proximity and the other based on similarity.

### 3.2.1 Seeded region growing

In this method, each seed pixel is given a label and added to a queue. While the queue is not empty, each pixel in the queue is removed and its neighbors are searched. If a neighbor,  $q$ , of a pixel labeled  $k$ , has not been labeled and the intensity of  $q$  is within a threshold of the intensity at the location of the seed  $k$ ,  $q$  is given the label  $k$  and  $q$  is added to the queue. The threshold used for growing is

$$t_k = w(s_k - I_m),$$

where  $s_k$  is the intensity at the location of the seed  $k$ ,  $I_m$  is the intensity minimum of the image, and  $w$  is a weight between  $[0,1]$ . This method gives a higher priority to proximity, since once a pixel is labeled it cannot be changed. A region is then created by the connected components of labeled pixels. Note that we could have used each marker region as a seed instead of the pixels from these regions. In this case, we could have used the mean intensity of the marker regions as the intensity of the seed pixel.

Figure 5 shows the results from seeded region growing applied to our sample images. Row 2 is obtained by using the foreground markers, without the cleanup step, as the seed pixels and in row 3, we use the pixels from the foreground markers obtained after the cleanup step described in Section 3.1.2. The regions shown are the ones which contain at least one foreground marker pixel which is simultaneously not a background marker pixel.

### 3.2.2 Seed competition

The second method of region growing, called seed competition,<sup>13</sup> grows a seed pixel until a neighboring pixel is reached that is more similar to another region. The growing process is similar to the seeded region growing algorithm except a priority queue is implemented, with priority given to neighboring pixels that are most similar to a labeled region. Two regions starting from different seeds are allowed to “compete” for a pixel. As a pixel can belong to either a coherent structure or the background, we use as seeds the pixels from both the background markers and the foreground markers obtained before the cleanup step. This results in a more accurately segmented image than using the foreground markers alone.

We start by adding each pixel from the seed regions to the priority queue which is emptied in the order of similarity. Each pixel in the queue is removed and its neighbors are searched. Suppose the neighbor,  $q$ , of pixel,  $p_s$ , is not labeled or it is labeled but  $q$  is more similar to  $p_s$  than the pixel,  $p_{s-i}$ , which decided its previous labeling. Then,  $q$  is given the same label as  $p_s$  and  $q$  is added to the priority queue based on its similarity to  $p_s$ . The similarity of  $p_s$  to  $q$  is measured using

$$\delta(p_s, q) = 1 - \exp\left(\frac{-1}{2\sigma^2}\left(\frac{I(p_s) + I(q)}{2} - I(s)\right)^2\right),$$

where  $\sigma$  is the standard deviation of the image,  $I(s)$  is the intensity at the seed  $s$ , and  $I(p_s)$  and  $I(q)$  are the intensities at  $p_s$  and  $q$ , respectively. Ties are broken by the first seed to reach the pixel. A region is then created by the connected components of labeled pixels. Sample results of the seed competition are presented in Figure 5, row 4. The regions shown are the ones which contain at least one foreground marker pixel which is simultaneously not a background marker pixel.

## 3.3 Model-Based Method Using 2-D Gaussians

Another approach to segmenting an image is to model the data. The ellipse-like shape of the coherent structures suggested the use of 2-D elliptic Gaussian functions. We implemented a modified version of the JMFIT algorithm in the Astronomical Image Processing System (AIPS).<sup>14,15</sup> JMFIT determines a least-square fit of a 2-D elliptic Gaussian to the data. The regions are specified by a center point  $c$  which is the location of the Gaussian mean, a

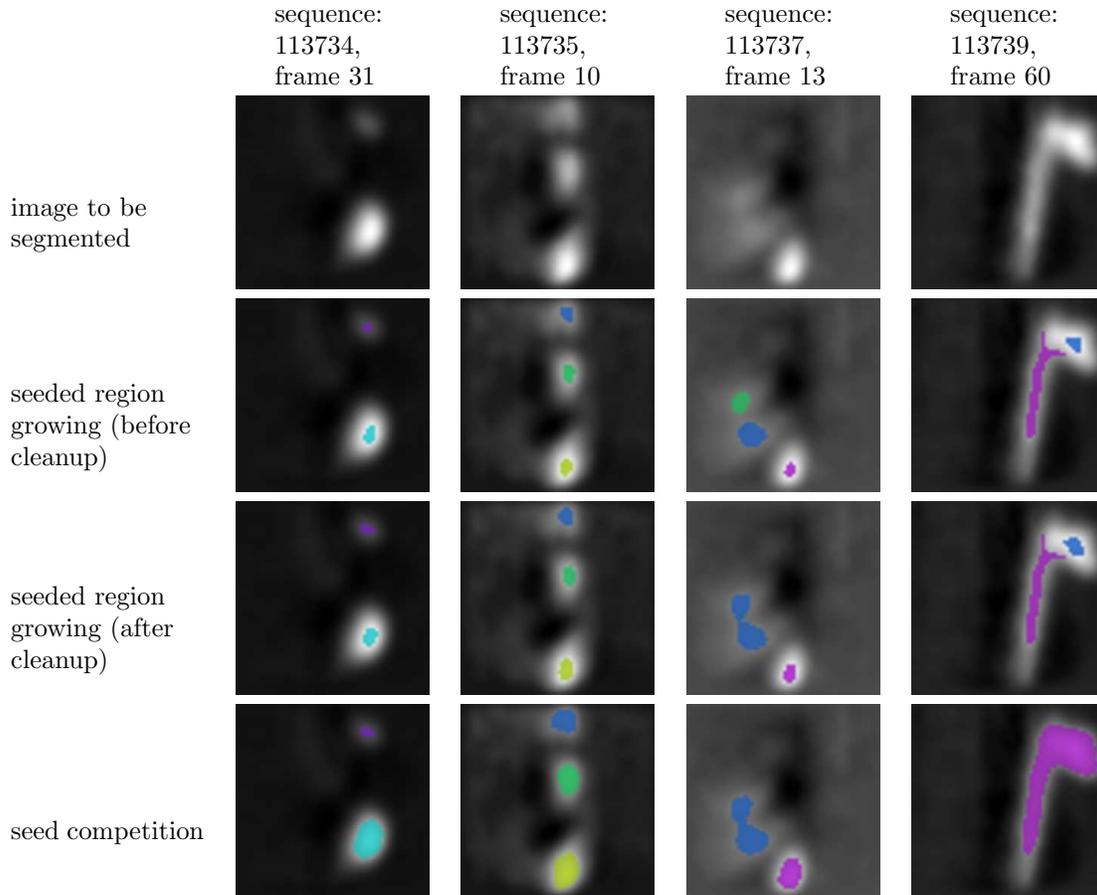


Figure 5. Results from the seeded region growing and seed competition methods: Row 1 - image to be segmented; row 2 - results of seeded region growing from foreground markers before cleanup (Figure 4 row 3); row 3 - results of the seeded region growing from foreground markers after cleanup (Figure 4 row 4); row 4 - results from the seed competition method.

major width  $w_{mj}$ , a minor width  $w_{mn}$ , and the angle  $\theta$  of the major axis counterclockwise from the north. As an initial guess, we use symmetric Gaussian functions ( $w_{mj} = w_{mn}$  and  $\theta = 0$ ) with center points at key foreground markers. These are a subset of the foreground marker pixels found before the cleanup step. We consider only those foreground marker pixels whose intensities are higher than all its neighbors within a 5 pixel radius. The initial guess for the major and minor width is found by searching radially from the foreground marker to the first value which is less than  $s_k/e$ , where  $s_k$  is the intensity at the location of the foreground marker  $k$  and  $e$  is the base of the natural logarithm.

Unfortunately, the number of Gaussian functions needed to best fit the data is unknown. Therefore, we begin with the number of key foreground markers as an initial guess. After the data is fitted, the key foreground markers are matched to the Gaussian functions. If unmatched key foreground markers exist, another fit is tried with the results from the previous fit and the unmatched key foreground markers forming the initial guesses. This continues for several iterations until convergence. To determine if a foreground marker is matched, the center of the Gaussian function must be within 5 pixels of the foreground marker. A Gaussian function may match with more than one foreground marker and a foreground marker may match multiple Gaussian functions; this is resolved by merging regions which are close in proximity and intensity. Figure 6, row 2 shows the estimate of the image reconstructed using all the Gaussian functions and in row 3, the matched ellipses. The ellipses shown in row 3 have a positive peak with a center not in the background. The axes of the ellipse are drawn at the full width at 0.75 of the maximum of the Gaussian functions. We observe that the estimate images appear similar to the images in row 1 and the ellipses overall coincide with the visible coherent structures.

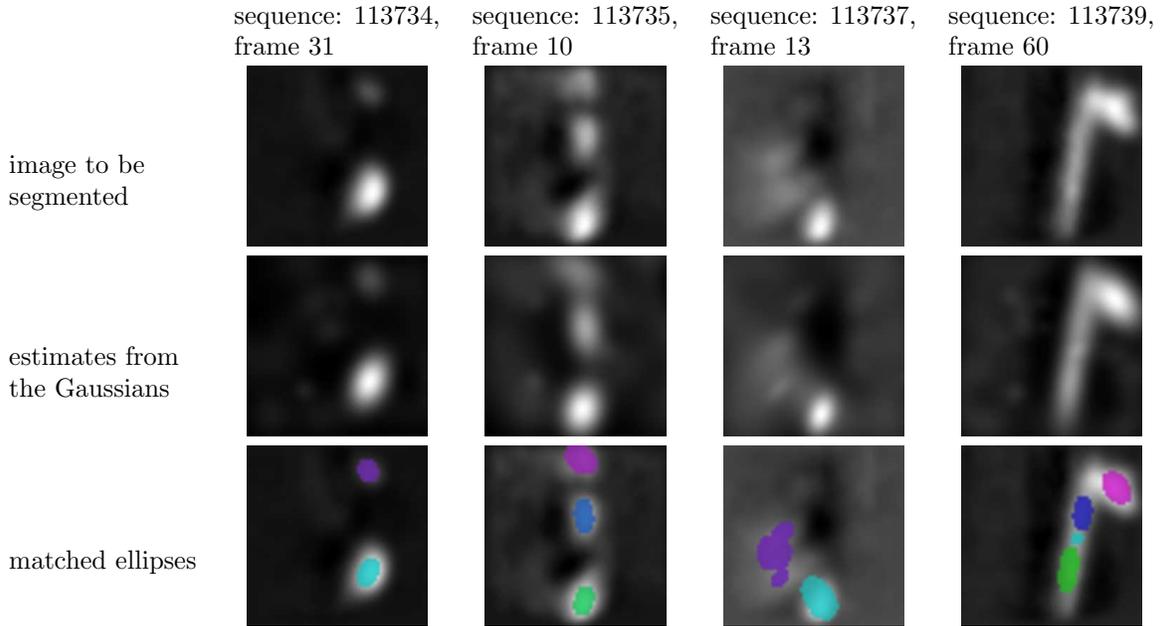


Figure 6. Results of segmentation by fitting 2-D elliptic Gaussians: Row 1 - images to be segmented; row 2 - estimated image reconstructed using the Gaussians and a maximum of 10 iterations; row 3 - ellipses with a positive peak and center points not in the background markers are shown highlighted on the original image.

#### 4. COMPARISON OF THE SEGMENTATION TECHNIQUES

In Section 3, we presented the results from different segmentation techniques on four sample images taken from four different NSTX sequences, each 300 frames long. As there is a wide variation in the coherent structures within a sequence, we applied these techniques to all frames of these four sequences. Based on a visual evaluation of the results, we make the following observations about the different segmentation methods:

- The basic immersion method is very simple and requires only a single parameter, namely, the immersion depth. However, care must be taken to interpret the resulting segmentation. If the immersion depth is specified as a fraction of the range of the intensities of the image, then we cannot directly compare the results of segmenting two consecutive frames from a sequence if the frames have very different intensity ranges. For example, Figure 7 shows two frames from the same sequence segmented using the same depth. The structures are correctly detected in frame 27, but appear to include a large part of the background in frame 59. In the latter image, as the structures are so diffuse, the results we obtain would be very sensitive to the immersion depth, regardless of how it is selected. If we specify the immersion depth as an absolute value, making it possible to compare results across frames, then we will need to set the threshold so that even the faint structures can be detected. However, this would result in merging a bright structure next to a faint structure or would pick up faint regions which meet the segmentation criterion, but which are not coherent structures.
- Any method which is based on the foreground and background markers is sensitive to both the existence and the location of the markers. There are several aspects to this which we discuss next. The first issue is that we consider as coherent structures only those regions which contain a foreground marker pixel which is not simultaneously a background marker pixel. As the structures are visually identified as bright blobs, we expect them to have a regional maximum, that is, a foreground marker. However, we also want to avoid those foreground markers which are also background markers. This constraint may miss structures of interest as shown in Figure 8 for the constrained watershed and seed competition methods. There are at least two regions, one each on either side of the central darker region, which should be included in the final segmentation, but are not as their foreground markers also lie in the background. Further, isolated

faint regions often get included in the background marker and therefore, are never detected as coherent structures until they become bright enough as in the faint blob near the top of Figure 9.

- Another problem arises due to the definition of the foreground marker which requires the marker pixels to be regional maxima. Consider Figure 9, which shows the results for seed competition and constrained watershed. The image contains a single bright blob with a fainter tail. While one may question if the tail should be a separate structure or part of the brighter blob, we first need to ensure that the tail is detected as a structure. However, this can be rather difficult. A single foreground marker is detected in the bright blob, but there is no marker in the fainter tail as no pixel in it satisfies the constraints of the maxima. Therefore, constrained watershed does not include the faint tail as a coherent structure. Seed competition, as well as seeded region growing, will not grow the bright region to include the tail as there is a large intensity difference between the two, with the fainter tail being more similar to the background. Such problems can also arise with watershed merging.
- The seeded region growing technique can give rise to strangely-shaped regions, regardless of whether the foreground markers are obtained before or after cleanup. If cleanup is used, the morphological operations flatten the low intensity regions, and the low intensity seeds grow into the background, extending beyond the structure of interest (see Figure 10). Even if the cleanup step is skipped as in Figure 11, we may get strangely shaped low-intensity regions. In addition, if the threshold for growing is set incorrectly, a bright region could stop growing prematurely. The two effects can combine to create a bright region which is separated from, but partially or completely surrounded by, a low intensity region, as in Figure 11.
- The results from any technique which depends on a threshold can be sensitive to the setting of the threshold. This includes the basic immersion method, watershed merging, and seeded region growing.
- The Gaussian fit method gives rather good estimates of the image using a combination of 2-D elliptic Gaussian functions. However, some of the Gaussians do not directly correspond to a coherent structure, though they are essential to getting a good estimate to the image. A challenge therefore is to translate the information in the Gaussians to the structures of interest in the image. Though we consider those Gaussians with a positive peak and a center not in the background, it is not clear if this is the best approach as seen in the last column of Figure 6.

## 5. CONCLUSIONS

In this paper we discussed some preliminary results from the application of image segmentation techniques to the analysis of experimental images of edge turbulence in nuclear fusion reactors. Using four image sequences from the gas puff imaging diagnostic on the National Spherical Torus Experiment, we showed how segmentation can be used to detect coherent structures in the images.

Our early experiences indicate that some of these techniques such as simple immersion and seeded region growing can be sensitive to the choice of thresholds, while others such as watershed merging and seed competition can depend on the choice of foreground and background markers. A key issue in many of the techniques is the determination of how we map the segmentation results to the coherent structures of interest. A complicating factor is that the structures are poorly understood empirically and not understood theoretically. In future, we plan to pursue the more promising techniques such as the 2-D Gaussian fit, seed competition, and constrained watershed to see if we can improve their performance by incorporating knowledge about the image into the choice of parameters.

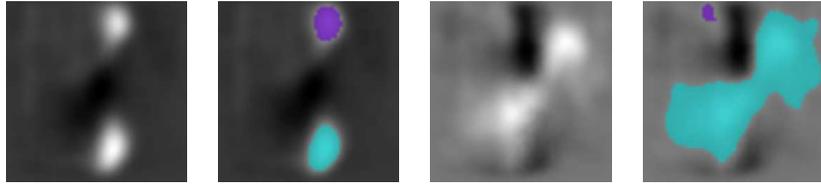


Figure 7. Original image and after segmentation with the basic immersion method (depth=0.5), sequence 113735, frames 27 and 59.

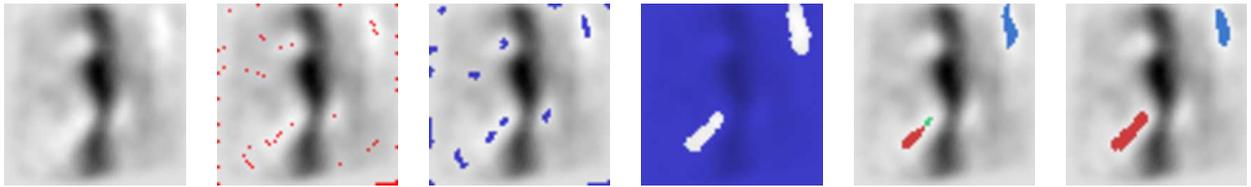


Figure 8. Original image, foreground markers before cleanup, foreground markers after cleanup, background markers, results from constrained watershed, and results from seed competition for sequence 113737, frame 69.

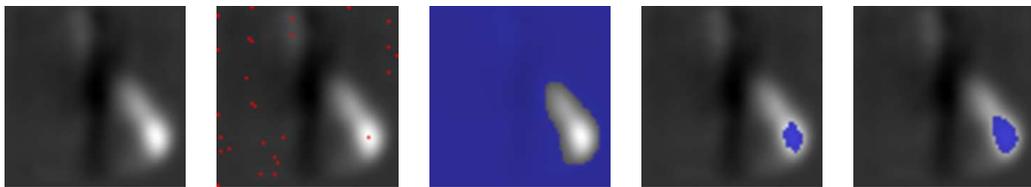


Figure 9. Original image, foreground markers before cleanup, background markers, results from constrained watershed, and results from seed competition for seq:113737, frame 179.

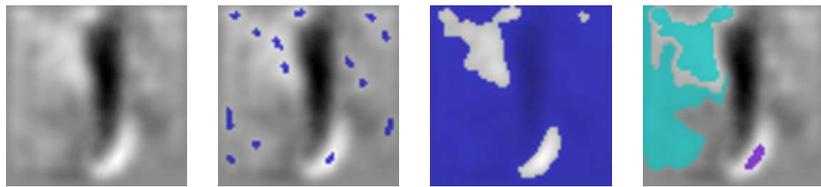


Figure 10. Original image, foreground markers after cleanup, background markers, and results from seeded region growing using foreground markers after cleanup for sequence 113734, frame 6.

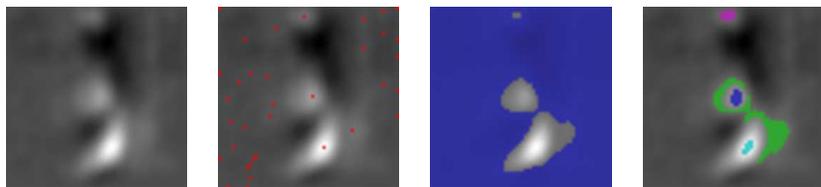


Figure 11. Original image, foreground markers before cleanup, background markers, and results from seeded region growing using foreground markers before cleanup for sequence 113734, frame 61.

## ACKNOWLEDGMENTS

We would like to thank Stewart Zweben of the Princeton Plasma Physics Laboratory (PPPL) for the image sequences and Stewart Zweben, Daren Stotler, and Ricardo Maqueda, all of PPPL, for the technical discussions. The software used in this work was developed by the members of the Sapphire scientific data mining project at LLNL.

A color version of this document can be found at <http://www.llnl.gov/CASC/sapphire>.

UCRL-CONF-233180: This work was performed under the auspices of the U.S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

## REFERENCES

1. ITER Web page, “The International Thermonuclear Experimental Reactor (ITER) project.” <http://www.iter.org/>, November 2004.
2. S. Zweben, “Search for coherent structure within tokamak plasma turbulence,” *Physics Fluids* **3**, pp. 974–982, 1985.
3. R. Maqueda, G. Wurden, D. Stotler, S. Zweben, B. LaBombard, J. Terry, J. Lowrance, V. Mastrocola, G. Renda, D. D’Ippolito, J. Myra, and N. Nishino, “Gas puff imaging of edge turbulence,” *Review of Scientific Instruments* **74**, pp. 2020–2026, March 2003.
4. D. D’Ippolito, J. Myra, S. Krasheninnikov, G. Yu, and A. Y. Pigarov, “Blob transport in the tokamak scrape-off-layer,” *Contributions to Plasma Physics* **44**, pp. 205–216, 2004.
5. S. Zweben, R. Maqueda, D. Stotler, A. Keesee, J. Boedo, C. Bush, S. Kaye, B. LeBlanc, J. Lowrance, V. Mastrocola, R. Maingi, N. Nishino, G. Renda, D. Swain, J. Wilgen, and the NSTX Team, “High-speed imaging of edge turbulence in NSTX,” *Nuclear Fusion* **44**, pp. 134–153, 2004.
6. S. Zweben, J. Boeda, O. Grulke, C. Hidalgo, B. LaBombard, R. Maqueda, P. Scarin, and J. Terry, “Edge turbulence measurements in toroidal fusion devices,” *Plasma Physics and Controlled Fusion* **49**, pp. S1–S23, 2007.
7. M. Sonka, V. Hlavac, and R. Boyle, *Image processing, analysis, and machine vision*, Brooks/Cole, San Francisco, CA, 2nd ed., 1999.
8. T. Munsat and S. Zweben, “Derivation of time-dependent 2-D velocity field maps for plasma turbulence studies,” *Review of Scientific Instruments* **77**, October 2006.
9. L. Vincent and P. Soille, “Watersheds in digital spaces: An efficient algorithm based on immersion simulations,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**, pp. 583–598, June 1991.
10. Mathworks, “Marker-controlled watershed segmentation.” <http://www.mathworks.com/>, 2006.
11. N. Otsu, “A thresholding selection method from grey-level histograms,” *IEEE Trans. on Systems, Man, and Cybernetics* **SMC-9**, pp. 62–66, January 1979.
12. K. Haris, S. Efstratiadis, and N. Maglaveras, “Watershed-based image segmentation with fast region merging,” *IEEE International Conference on Image Processing* **3**, pp. 338–342, 1998.
13. A. X. Falcao, P. Miranda, A. Rocha, and F. Bergo, “Object detection by k-connected seed competition,” in *Proceedings of the XVII Brazilian Symposium on Computer Graphics and Image Processing*, pp. 97–104, October 2005.
14. National Radio Astronomy Observatory, “Astronomical Image Processing System (AIPS) cookbook.” <http://www.aoc.nrao.edu/aips/>, December 2006.
15. R. L. White, R. H. Becker, D. J. Helf, and M. D. Gregg, “A catalog of 1.4 GHz radio sources from the FIRST survey,” *The Astrophysical Journal* **475**(1), pp. 479–493, 1997.